arXiv:1703.07419v1 [cs.NI] 21 Mar 2017

Optimal Routing for Delay-Sensitive Traffic in Overlay Networks

Rahul Singh and Eytan Modiano

Abstract

We design dynamic routing policies for an overlay network which meet delay requirements of real-time traffic being served on top of an underlying legacy network, where the overlay nodes do not know the underlay characteristics. We pose the problem as a constrained MDP, and show that when the underlay implements static policies such as FIFO with randomized routing, then a decentralized policy, that can be computed efficiently in a distributed fashion, is optimal. When the underlay implements a dynamic policy, we show that the same decentralized policy performs as well as any decentralized policy. Our algorithm utilizes multi-timescale stochastic approximation techniques, and its convergence relies on the fact that the recursions asymptotically track a nonlinear differential equation, namely the replicator equation.

I. Introduction

Overlay networks are a novel concept to bridge the gap between what modern Internet-based services need and what the existing networks actually provide, and hence overcome the shortcomings of the Internet architecture [1], [2]. The overlay creates a virtual network over the existing underlay, utilizes the functional primitives of the underlay, and supports the requirements of modern Internet-based services which the underlay is not able to do on its own (see Fig. 1). The overlay approach enables new services at incremental deployment cost.

The focus of this paper is on developing efficient overlay routing algorithms for data which is generated in real-time, e.g., Internet-based applications such as banking, gaming, shopping, or live streaming. Such applications are sensitive to the end-to-end delays experienced by the data packets.

Dynamic routing policies for multihop networks have been traditionally studied in the context where all nodes are controllable [3], [4]. Our setup, however, allows only a subset of the network nodes (overlay) to make dynamic routing decisions, while the nodes of the *legacy network* (underlay) implement simple policies such as FIFO combined with fixed path routing. This approach introduces new challenges because the overlay nodes do not have knowledge of the underlay's topology, routing scheme or link capacities. Thus, the overlay nodes have to *learn* the optimal routing policy in an *online* fashion which involves an exploration-exploitation trade-off between finding low delay paths and utilizing them. Moreover, since the network conditions and traffic demands may be time-varying, this also involves consistently "tracking" the optimal policy.

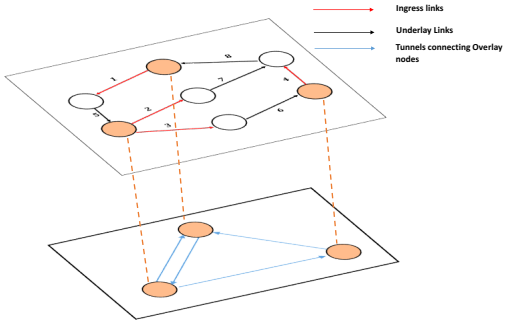

Fig. 1. The bottom plane shows the complete network in which the Overlay nodes are colored. The top plane visualizes the overlay network, in which a tunnel corresponds to an existing, possibly randomized underlay path. The ingress links $1, 2, 3, 4$ inject traffic from overlay nodes into the underlay nodes.

A. Previous Works

The use of overlay architecture was originally proposed in [5] to achieve network resilience by finding new paths in the event of outages in the underlay. In a related work [6] considered the problem of placing the underlay nodes in an optimal fashion in order to attain the maximum "path diversity". In [7] the authors consider optimal overlay node placement to maximize throughput, while [8] develops throughput optimal overlay routing in a restricted setting.

While many works [9], [10] use end-to-end feedbacks for delay optimal routing, these works ignore the *queueing aspect*, and hence the delay of a packet is assumed to be independent of the congestion in the underlay network. Early works on delay minimization [11], [12] concentrated on quasi-static routing, and do not take the network state into account while making routing decisions. The existing results on dynamic routing explicitly assume that all the network nodes are controllable, and typically analyze the performance of algorithms when the network is heavily loaded [4]. However, in the heavy traffic regime, the delays incurred under any policy are necessarily large, and thus not suitable for routing of real-time traffic. Finally, we note that the popular backpressure algorithm is known to perform poorly with respect to average delays [13].

B. Contributions

In contrast to the approaches mentioned above, we consider a network where only a subset of nodes (overlay) are controllable, and propose algorithms that meet average end-to-end

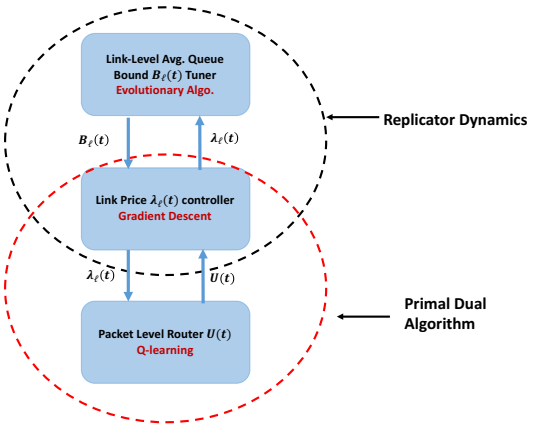

Fig. 2. The proposed 3 layer Price-based Overlay Controller (POC) comprising of (from bottom to top) i) Overlay nodes making packet-level routing decisions $U(t)$, ii) link-level price controller which manipulates the prices $\lambda_\ell(t)$, iii) Link-level average queue threshold manipulator which tunes the $B_\ell(t)$. The interactions between the top 2 layers are described by a nonlinear ode called replicator dynamics, while the bottom 2 layers constitute a primal-dual algorithm.

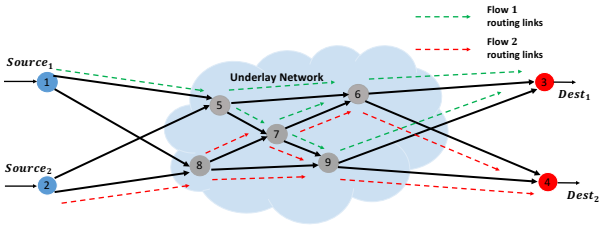

Fig. 3. An overlay network comprising of 2 source destination pairs $(1,3)$ and $(2,4)$. Since the routes taken by packets sent at ingress links $(1,5)$ and $(2,8)$ share common underlay links $(7,9)$ and $(7,6)$, an increase in traffic intensity sent on either of the ingress links may also increase the delay suffered by packets sent on the other ingress link.

delay requirements imposed by applications. Our algorithms are decentralized and perform optimally irrespective of the network load.

It follows from Little's law [14] that for a stable network, the objective of meeting an average end-to-end delay requirement can be replaced by keeping the average queue lengths below some value B. In this work, the problem of maintaining the average queue lengths below B is divided into two sub-problems: i) distributing the bound B into link-level average queue thresholds B_ℓ such that these link-level bounds can be satisfied under some policy, ii) designing an overlay routing policy that meets the link-level queue bounds imposed by i).

We obtain an efficient decentralized solution to ii) by introducing the notion of "link prices" that are charged by links. The link prices induce cooperation amongst the overlay nodes, thereby producing decentralized optimal routing policy, and are in spirit with the Kelly decomposition method for network utility maximization [15]. The average queue lengths are adaptively controlled [16], [17] by manipulating the link prices, but unlike previous works which utilize Kelly decomposition in a static deterministic setting [18], we perform a stochastic dynamic optimization with respect to the routing decisions.

In order to solve i) we provide an adaptive scheme which follows the replicator dynamics [19]. Finally, the solutions to i) and ii) are combined to yield a 3 layered queue control overlay routing scheme, see Fig. 2.

Our problem also has close connections to the restless Multi Armed Bandit problem (MABP) [20], [21]. Our scheme takes an "explore-exploit strategy" in absence of knowledge regarding the underlay network's characteristics, and learns the optimal routing policy in an online fashion using the data obtained from network operation. The routing decisions on the various ingress links (see Fig 3) correspond to the bandit "arms", while the average end-to-end network delays are the *unknown* rewards. Since the packets injected by different ingress links share common underlay links on their path, routing decisions at an ingress link ℓ affect the delays incurred by packets sent on different ingress link $\hat{\ell}$ (see Fig. 3). This introduces dependencies amongst the Bandit arms, and hence the decision space grows exponentially with the number of ingress overlay links. Consequently we cannot apply existing MABP algorithms, and must develop simpler algorithms that suit our needs.

Our goal is to develop decentralized policies to control the end-to-end delays, which can be computed efficiently in a parallel and distributed fashion [22]. We show that if the underlay implements a simple static policy, then there exists a decentralized policy that is optimal. When the underlay is allowed to use dynamic policies, we provide theoretical guarantees for our decentralized policies.

We begin in Section II by describing the set-up, and pose the problem of designing an overlay policy to keep the average end-to-end delays within a pre-specified bound, as a constrained Markov Decision Process (CMDP) [23]. Section III-A1 solves the problem of meeting link-level average queue bounds. It is shown that the routing decisions across the flows can be decoupled if the links are allowed to charge prices for their usage. This *flow-level decomposition* technique significantly simplifies the policy design procedure, and also ensures that the resulting scheme is decentralized. Section IV employs an evolutionary algorithm to tune the link-level average queue bounds, and proves the convergence properties. Section VI discusses several useful extensions. Section **??** compares the performance of our proposed algorithms with the existing algorithms.

II. Problem Formulation

We will first describe the system model, and then proceed to pose the problem of bounding the average end-to-end delays.

A. System Model

The network is represented as a graph $G = (V, E)$, where V is the set of nodes, and E is the set of links. The network evolves over discrete time-slots. A link $\ell = (i, j) \in E$ with

capacity $C_l(t), t = 1, 2, \ldots$ implies that node i can transmit C_l packets to node j at time t. We allow for the link capacities $C_\ell(t)$ to be stochastic, i.e., $C_\ell(t)$ depends on the state of link ℓ at time t. We will assume that the link states are i.i.d. across time [1].

Multiple flows $f = 1, 2, \ldots, F$ share the network. Each flow f will be associated with a source node s_f and destination node d_f. We will assume that the packet arrivals at each source node s_f are i.i.d. across time and flows[2]. Mean arrival rate at source node s_f will be denoted by A_f. The number of arrivals at any source node are uniformly bounded across time and flows.

There are two types of nodes in G: i) *overlay*: Those that can make *dynamic* routing and scheduling decisions based on the network state, ii) *underlay*: Those that implement FIFO scheduling combined with randomized routing, on a per flow basis.

The subgraph induced by the underlay nodes will be called the underlay network or just underlay. In order to make the exposition simpler and avoid cumbersome notation, we will make some simplifying assumptions[3]. Firstly, the overlay will be assumed to be composed entirely of source and destination nodes. Thus the set of overlay nodes is given by,

$$\{i \in V : i = s_f \text{ or } i = d_f \text{ for some flow } f = 1, 2, \ldots, F\}.$$

Throughout the paper we will only consider the case where the underlay is weakly connected, i.e., there are no multiple alternating segments of overlay-underlay connecting the source-destination pairs. We will assume that there are no links connecting two overlay nodes directly, i.e., the overlay nodes are connected in the overlay network only through underlay tunnels (see Fig. 1). Also, the flows do not share source nodes.

Ingress Links : The network links $\{\ell = (i, j) : i \in \text{Overlay}, j \in \text{Underlay}\}$ will be referred to as the ingress links. These are the outgoing links from overlay nodes that connect with the underlay, and hence these are precisely the links where the overlay routing decision take place (see Fig. 1).

Tunnel: For each outgoing ingress link ℓ from s_f, and a destination node i belonging to overlay, we say that the tunnel $\tau_{\ell,i}$ connects the node s_f to i in the overlay network, see Fig. 1. We note that the actual path taken by a packet sent on an ingress link, which is a sequence of underlay links, depends on the randomized routing decision taken by the underlay, and is not known at the overlay.

Underlay Operation : The network evolves over discrete time slots indexed $t = 0, 1, \ldots$. Each *underlay link* ℓ maintains a separate queueing buffer for each flow in which it stores packets belonging to that flow (see Fig. 4). An underlay link ℓ is shared by queues belonging to flows whose routes utilize link ℓ. In order to simplify the exposition of ideas, we begin

[1]Our analysis extends in a straightforward manner for the case when the states evolve as a finite state Markov process.

[2]Our analysis can easily be extended for the case when arrivals are governed by a finite state Markov process.

[3]Our algorithms, and their analysis can be easily extended to the case where these assumptions are not satisfied. We choose to present the simple case in order to simplify the exposition of ideas, and avoid unnecessary notation.

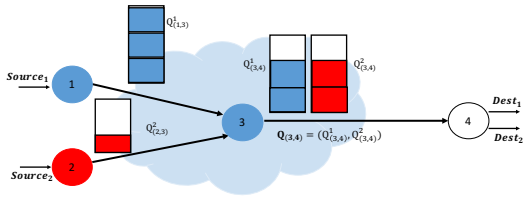

Fig. 4. An overlay network comprising of 2 flows that share the same destination node 4. Each link ℓ maintains separate queues for the flows f that are routed on it. Queues on link $(3, 4)$ are given by the vector $\boldsymbol{Q}_{(3,4)} = \left(Q^1_{(3,4)}, Q^2_{(3,4)}\right)$.

by assuming that each underlay link ℓ implements a static scheduling policy by sharing the available capacity at any time t amongst the flows in some pre-decided ratio. Thus, there are constants $\mu_{\ell,f}, f = 1, 2, \ldots, F \quad \ell \in E$ satisfying $\mu_{\ell,f} \geq 0, \sum_f \mu_{\ell,f} = 1$. At each time t, flow f receives a proportion $\mu_{\ell,f}$ of the available link capacity at link ℓ. Note that in case a flow f does not use the link ℓ for routing, then $\mu_{\ell,f} = 0$. The links can use randomization in order to allocate the capacity in the ratio $\mu_{\ell,f}$.

The underlay employs a randomized routing discipline. Thus, once a flow f packet is successfully transmitted on an underlay link ℓ, it is routed with a probability $P^f(\ell, \hat{\ell})$ to link $\hat{\ell}$. Note that $\sum_{\hat{\ell}} P^f(\ell, \hat{\ell}) = 1 \quad \forall f, \ell$, and the routing probabilities are allowed to be flow-dependent. We denote by R_f, the set of links ℓ which are used for routing flow f packets. Note that this includes static routing such as shortest path.

Our analysis extends easily to include the case where the underlay can utilize dynamic routing-scheduling disciplines such as longest queue first, priority based scheduling, or even the Backpressure policy [3]. We discuss the extension in Section VI.

Information available to the Overlay: Overlay nodes do not know the underlay's topology, probability laws governing the random link capacities, nor the policy implemented by it. In Sections III-A1, and IV we will assume that each source node s_f knows the underlay queue lengths corresponding to its flow. Later, we will relax this assumption and assume that source nodes know only the end-to-end delays corresponding to its flows.

Actions available to the Overlay: Let $\boldsymbol{Q}(t) := (\boldsymbol{Q}_1(t), \boldsymbol{Q}_2(t), \ldots, \boldsymbol{Q}_L(t))$, where each $\boldsymbol{Q}_\ell(t)$ is a vector containing queue lengths $Q^f_\ell(t)$ of all the flows f that are routed through link ℓ (see Fig. 4). At each time $t \geq 0$, each overlay source node s_f decides its action $\boldsymbol{U}^f(t) := \{U^f_\ell(t)\}_{\ell=(s_f,i)}$,where $U^f_\ell(t)$ is the number of flow f packets sent to the queueing buffer of the ingress link ℓ at time t. This decision is made based on the information available to it until time t which includes $\{\boldsymbol{U}^f(s), \boldsymbol{Q}^f(s), \lambda(s)\}_{s=0}^{t-1}$, where $\boldsymbol{Q}^f(t)$ is the vector of queue lengths corresponding to the flow f, and $\lambda(s) = \{\lambda_\ell(s)\}_{\ell \in E}$ is the vector comprising of link prices at time s.

Scheduling Policy: We will denote vectors in boldface. A routing policy π is a map from $\boldsymbol{Q}(t)$, the vector of queue lengths at time t, to a scheduling action $\boldsymbol{U}(t) = \{\boldsymbol{U}^f(t)\}_{f=1,2,\ldots,F}$ at time t. The action $\boldsymbol{U}^f(t)$ determines the number of packets to be routed at time t on each of the ingress links belonging to flow f.

Policy at Underlay Throughout, we will assume that the underlay links implement a simple static policy in which each link ℓ shares the link capacity C_ℓ amongst the flows in some constant ratio, even if some of the queues $Q_\ell^f(t)$ are empty. In Section VI we will provide guidelines to consider the case when such an assumption is not true, i.e., the underlay network implements a more complex policy.

B. Objective: Keep Average Delays Bounded by B

We now pose the problem of designing an overlay routing policy that meets a specified requirement on the average end-to-end delays incurred by packets as a constrained MDP [23].

For a stable network, it follows from the Little's law [14], that for a fixed value of the mean arrival rate, the mean delay is proportional to the average queue lengths. Thus, we will focus on controlling the average queue lengths instead of bounding the end-to-end delays. A bound on the average queue lengths will imply a bound on the average end-to-end delays.

Since now our objective is to control the average queue lengths, the state ([24]) of the network at time t is specified by the vector $\boldsymbol{Q}(t)$. Also, we let $\|\boldsymbol{Q}(t)\| := \sum_{\ell,f} Q_\ell^f(t)$ denote the total queue length, i.e. the 1 norm of the vector $\boldsymbol{Q}(t)$. We also let $\boldsymbol{Q}^f(t) := \{Q_\ell^f(t)\}_{\ell \in R_f}$ denote the vector containing queue lengths belonging to flow f. The queue dynamics are given by,

$$Q_\ell^f(t+1) = \left(Q_\ell^f(t) - D_\ell^f(t)\right)^+ + A_\ell^f(t), \forall f, \ell,$$

where $D_\ell^f(t), A_\ell^f(t)$ are the number of flow f departures and arrivals respectively at link ℓ at time t. The arrivals could be external, or due to routing after service completions at other links.

If the overlay knew the underlay topology, link capacities and routing policy, the overlay could solve the following constrained MDP [23] in order to keep the average queue lengths less than the threshold B,

$$\min_{\pi} 0 \tag{1}$$

$$\lim_{T\to\infty} \frac{1}{T}\mathbb{E}_\pi \left\{\sum_{t=1}^{T} \|\boldsymbol{Q}(t)\|\right\} \leq B, \tag{2}$$

where expectation is taking with respect to the (possibly random) routing policies at the overlay and underlay, and the randomness of arrivals and link capacities. If we assume that the network queue lengths are uniformly bounded, i.e., $Q_\ell^f(t) \leq C, \quad \forall \ell \in E, f = 1, 2, \ldots, F$, then standard theory from constrained MDP [23] tells us that the above is solved by a stationary randomized policy. Thus $\boldsymbol{U}(t) = \{\boldsymbol{U}^f(t)\}_{f=1,2,\ldots,F}$, the overlay routing decision at time t is solely determined by the network state, i.e., $\boldsymbol{U}(t) = h(\boldsymbol{Q}(t))$.

The assumption of bounded queues is not overly restrictive since the links can simply drop a packet if their buffer overflows[4].

We note that we can also consider the set-up in which each flow f has an average end-to-end delay requirement, and the objective is to design an overlay routing policy that meets the requirements imposed by each flow f. We exclude this set-up to make the presentation easier, and instead focus only on the case where the delays summed up over flows is to be kept below a certain threshold. Details regarding this extension can be found in Section VI.

Henceforth we will write $\|\bar{\boldsymbol{Q}}\|_\pi$ to denote the total steady state average queue lengths under the application of policy π. We will suppress the dependency on π and use $\|\bar{\boldsymbol{Q}}\|$ instead of $\|\bar{\boldsymbol{Q}}\|_\pi$. Similarly for $\|\bar{\boldsymbol{Q}}_\ell\|, \bar{Q}_\ell^f$ etc.

III. Optimal Policy Design

We now begin our analysis when the underlay implements a static policy. We will show that it is possible to construct an optimal decentralized overlay routing policy using a 3 layered controller as shown in Figure 2. The topmost layer will manipulate "link-level average queue thresholds" $B_\ell(t)$ based on the link prices $\lambda_\ell(t)$. The link prices $\lambda_\ell(t)$ are representative of the instantaneous congestion at link ℓ. The next two layers, link-price controller, and packet-level decision maker will collectively try to meet the bounds $B_\ell(t)$ by using a primal dual algorithm described below. The packet-level decision maker will utilize the link prices $\lambda_\ell(t)$ in order to make routing decisions $\boldsymbol{U}(t)$. It will transmit packets on routes which utilize links with lower prices. The link-price controller will then observe the congestion that results from the routing decisions $\boldsymbol{U}(t)$, and manipulate the prices accordingly in order to direct more traffic towards links on which the average queues are less than the threshold $B_\ell(t)$. The interaction between the top two layers is described by a nonlinear ordinary differential equation (ode) called the replicator equation [19]. Fig. 2 depicts the interactions between these layers. We call our 3 layered adaptive routing policy Price-based Overlay Controller (POC). We now develop these algorithms in a bottom to top fashion.

A. Link-level Design

We notice that in order to satisfy the constraint $\sum_\ell \|\bar{\boldsymbol{Q}}_\ell\| \leq B$ imposed in the CMDP (1), a policy π needs to appropriately coordinate the individual link-level average queue lengths $\|\bar{\boldsymbol{Q}}_\ell\|$ in order that their combined value is less than B. Such a task is difficult because the link-level average queues $\|\bar{\boldsymbol{Q}}_\ell\|$ have complex interdependencies between them that are described by the *unknown* underlay network structure.

[4]The steady state packet loss probabilities can be made arbitrarily small by choosing the bound C to be sufficiently large. Moreover, such an assumption automatically guarantees stability of queues, and lets us focus on our key objective of controlling network delays.

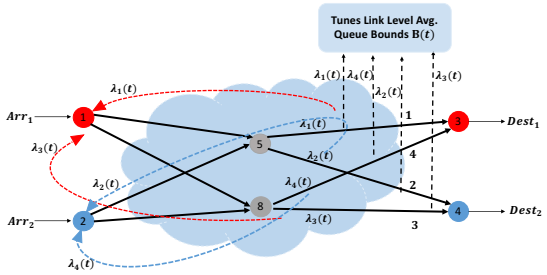

Fig. 5. An overlay network comprising of 2 source destination pairs $(1,3)$ and $(2,4)$. The introduction of link prices $\lambda_\ell(t)$ decouples the routing decisions at the overlay and results in a decentralized scheme. It also serves as a mediator between the overlay nodes making routing decisions, and the tuner which sets the queue thresholds B_ℓ.

In view of the above discussion, we begin by considering a simpler problem, one in which the tolerable cumulative average queue bound B has already been divided into link-level "components" $\{B_\ell\}_{\ell\in E}$ that satisfy $\sum_{\ell\in E} B_\ell = B$, and the task of overlay is to keep the average queues on each link ℓ bounded by the quantity B_ℓ. An efficient scheme should be able to meet the link-level bounds B_ℓ in case they are achievable under *some* routing scheme.

Thus, we consider the following Constrained MDP [23],

$$\min_{\pi} \quad 0 \tag{3}$$

$$\text{s.t. } \|\bar{\boldsymbol{Q}}_\ell\| \le B_\ell, \quad \forall \ell \in E, \tag{4}$$

where in the above, expectation is taken with respect to the overlay policy π, randomness of packet arrivals, network link capacities and the underlay policy. Next, we derive an iterative scheme that yields a decentralized policy which solves the above CMDP, i.e., satisfies the link-level average queue length bounds in case they are achievable under some policy. The task of choosing the bounds B_ℓ will be deferred until Section IV.

1) Flow Level Policy Decomposition: We now show that the problem (3) of routing packets under average delay constraints at each link $\ell \in E$ admits a decentralized solution of the following form. Each underlay network link ℓ charges a "holding cost" at the rate of λ_ℓ units per unit queue length Q_ℓ^f from each flow f that utilizes link ℓ. The holding charge collected by link ℓ from flow f at time t is thus $\lambda_\ell(t)Q_\ell^f(t)$. The flows are provided with the link prices $\{\lambda_\ell\}_{\ell\in E}$, and they schedule their packet transmissions in order to minimize their average holding costs $\sum_\ell \lambda_\ell \bar{Q}_\ell^f$. Under such a scheme, the prices λ_ℓ intend to keep the network traffic from "flooding" the links, thus enabling them to meet the average queue threshold of B_ℓ units.

The routing scheme described above is *decentralized*, meaning that each flow f makes the routing decisions solely based on its individual queue lengths $\boldsymbol{Q}^f(t)$, and the link prices λ, and not on the basis of global state vector $\boldsymbol{Q}(t)$. Thus, the problem of routing packets in order to meet the average queue thresholds B_ℓ neatly decomposes into F sub-problems, one for each flow. The sub-problem for flow f is to route its packets in order to minimize its individual holding cost $\sum_\ell \lambda_\ell \bar{Q}_\ell^f$.

We now proceed to solve the CMDP (3). Throughout this section we will assume that (3) is feasible, i.e., there exists a routing policy π under which the average queue lengths are less than or equal to the thresholds B_ℓ. Let $\lambda_\ell \ge 0$ denote the Lagrange multiplier associated with the constraint $\|\bar{\boldsymbol{Q}}_\ell\| \le B_\ell$, and let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_L)$. If π is a stationary randomized policy ([23], [25]), then the Lagrangian for the problem (3) is given by [23],

$$\begin{aligned}\mathcal{L}(\pi,\lambda) &= \sum_\ell \lambda_\ell \left(\|\bar{\boldsymbol{Q}}_\ell\| - B_\ell\right) \\ &= \sum_\ell \lambda_\ell \left(\sum_f \bar{Q}_\ell^f\right) - \sum_\ell \lambda_\ell B_\ell \\ &= \sum_f \left(\sum_{\ell\in R_f} \lambda_\ell \bar{Q}_\ell^f\right) - \sum_\ell \lambda_\ell B_\ell, \end{aligned} \tag{5}$$

where the second equality follows from the relation $\|\boldsymbol{Q}_\ell(t)\| = \sum_f Q_\ell^f(t)$. The Lagrangian thus decomposes into the sum of flow f "holding costs" $\sum_{\ell\in R_f} \lambda_\ell \bar{Q}_\ell^f$, where R_f is the set of links utilized for routing packets of flow f.

We now make the following important observation. Since the links share the service amongst its queues in a manner irrespective of the queue length $\boldsymbol{Q}_\ell(t)$, the service received by flow f on link ℓ at any time t, is determined only by the capacity of link ℓ at time t, and queue length $Q_\ell^f(t)$. Thus, the evolution of queue length process $\boldsymbol{Q}^f(t), t \ge 0$ is completely described by $\boldsymbol{U}^f(t), t \ge 0$, i.e., the routing actions chosen for flow f at source node s_f, and the probability laws governing capacities of links ℓ that are used to route flow f packets. Let us denote by π_f a policy for flow f that maps $\boldsymbol{Q}^f(t)$ to $\boldsymbol{U}^f(t)$. Since for a fixed value of λ, the cost $\sum_{\ell\in R_f} \lambda_\ell \bar{Q}_\ell^f$ incurred by each flow f can be optimized independently of the policies chosen for other flows, it then follows from (5) that,

Lemma 1: The dual function corresponding to the CMDP (3) is given by,

$$\begin{aligned} D(\lambda) &= \min_\pi \mathcal{L}(\pi,\lambda) \\ &= \sum_f \min_{\pi_f} \sum_{\ell\in R_f} \lambda_\ell \bar{Q}_\ell^f - \sum_\ell \lambda_\ell B_\ell., \end{aligned} \tag{6}$$

and hence the policy π which minimizes the Lagrangian for a fixed value of λ is *decentralized*.

We now provide a precise definition of a decentralized policy.

Decentralized Policy: A policy π is said to be decentralized if $\boldsymbol{U}^f(t)$, i.e., the overlay decisions regarding packets of flow f, are made based solely on the knowledge of flow f queues $\boldsymbol{Q}^f(t)$, i.e., $\boldsymbol{U}^f(t) = h_f(\boldsymbol{Q}^f(t))$. If π is decentralized, we let $\pi = \otimes_f \pi_f$, i.e., π is uniquely identified by describing the policies π_f for each individual flow f. Upon associating a policy with the steady state measure that it induces on the joint state-action pairs, we see that the decentralized policy corresponds to the product measure $\otimes_f \pi_f$ on the space

$\otimes_f(Q_f, U_f)$, where Q_f is the space on which $\boldsymbol{Q}^f(t)$ resides, and U_f is the action space corresponding to the routing decisions for flow f.

Let $\lambda^\star$ be the value of the Lagrange multiplier that solves the dual problem, i.e.,

$$\lambda^\star \in \arg\max_{\lambda \geq \mathbf{0}} D(\lambda), \tag{7}$$

where $\mathbf{0}$ is the $|E|$ dimensional vector all of whose entries are 0.

The average costs incurred under a policy π can be considered as a dot product between the steady state probabilities induced over the joint state-action pair under π, and the one-step cost under state-action pairs. Hence, a CMDP can equivalently be posed as a linear program [26], and is convex. Thus, the duality gap corresponding to the problem (3), and its dual (7) is zero [27]. Let us denote by $\pi_f^\star(\lambda)$ the policy for flow f which minimizes the holding cost $\sum_{\ell \in R_f} \lambda_\ell \bar{Q}_\ell^f$.

It then follows from the above discussion that,

Theorem 1: There exists $\lambda^\star = \{\lambda_\ell^\star\} \geq \mathbf{0}$ such that the CMDP (3) is solved by the policy $\pi^\star$ which implements for each flow f the corresponding policy $\pi_f^\star(\lambda^\star)$. Thus, $\pi^\star = \otimes_f \pi_f^\star(\lambda^\star)$, i.e., the optimal policy for the CMDP (3) is decentralized.

The consideration of the dual problem corresponding to CMDP (3) greatly reduces the problem complexity. We were able to show that the optimal policy is decentralized. As will be discussed in Section VIII-B, this has several advantages. However, the following two issues need to be addressed in order to compute the optimal policy $\pi^\star$.

1) Computation of the policies $\pi_f^\star(\lambda)$ requires the knowledge of underlay topology, distribution of underlay links' capacities and statistics of packet arrival processes. These are, however, not known at overlay. Moreover, in case the parameters are time-varying, the policy $\pi_f^\star(\lambda)$ necessarily needs to adapt to the changes.
2) Optimal policy $\pi^\star$ in Theorem 1 needs to know the vector of optimal "link prices" $\lambda^\star$ that solve the dual problem (7).

Next, we resolve these issues.

IV. Adaptively Tuning average Queue bounds B_ℓ

The scheme (23)-(24) developed in Section III-A1 keeps the average delays bounded by thresholds $\{B_\ell\}_{\ell \in E}$ that have been provided by the network operator. However, it works only if the thresholds $\{B_\ell\}_{\ell \in E}$ are achievable. Thus, while developing the scheme we had implicitly assumed that the overlay knew that the link-level average queue threshold profile $\{B_\ell\}_{\ell \in E}$ could be met, and thereafter it could utilize the Algorithm of Theorem 4 in order to meet the delay requirements.

But, the assumption that the overlay is able to characterize the set of achievable thresholds $\{B_\ell\}_{\ell \in E}$ is hard to justify in practice. Characterization of the set of achievable $\{B_\ell\}_{\ell \in E}$ is difficult because of the complex dependencies between various link-level delays. Furthermore, in order to calculate the average delay performance for any fixed policy π, the overlay needs to know the underlay characteristics, which we assume are unknown to the overlay. Thus, in this section, we devise an adaptive scheme that measures the link prices λ_ℓ that result when the scheme (23)-(24) is applied to the network, and utilizes them to iteratively tune the delay thresholds B_ℓ until an achievable vector $\boldsymbol{B}$ satisfying $\sum_\ell B_\ell = B$ is obtained. The key feature of the resulting scheme *is that it does not require the knowledge of the underlay network characteristics*. It increases the bounds B_ℓ for links with "higher prices", and decreases B_ℓ for links with "lower prices", while simultaneously ensuring that the new bounds sum up to B.

It is shown in Lemma 9 that the price iterations in (23) converge for each fixed $\boldsymbol{B}$. Let us denote by $\bar{\lambda}(\boldsymbol{B}) := \{\bar{\lambda}_\ell(\boldsymbol{B})\}_{\ell \in E}$ the vector comprising of link prices that result when the average queue thresholds are set at $\boldsymbol{B}$, and the iterations (23)-(24) are performed untill convergence. Thus, $\bar{\lambda}(\boldsymbol{B})$ denotes the steady-state link prices when the scheme of Theorem 4 is utilized with thresholds set to $\boldsymbol{B}$. We propose the following iterative scheme to tune the delay budget vector $\boldsymbol{B}(t)$,

$$\begin{aligned} &B_\ell(k+1) \\ &= \Gamma\left(B_\ell(k) + \gamma_k \left\{ B_\ell(k) \left(\bar{\lambda}(\boldsymbol{B}(k)) - \sum_\ell \bar{\lambda}_\ell(\boldsymbol{B}(k)) B_\ell(k) \right) \right\}\right), \\ &\quad \forall \ell \in E, \end{aligned} \tag{8}$$

where k is the iteration index, and $\Gamma(\cdot)$ is the projection operator that maps the vector $\boldsymbol{B}(k) \in \mathbb{R}^L$ onto the L dimensional simplex

$$S = \left\{ \boldsymbol{x} : \sum_{i=1}^{L} x_i = B, \quad x_i \geq 0 \quad \forall i = 1, 2, \ldots, L \right\}.$$

Projection operator $\Gamma(\cdot)$ is required in order that the iterates satisfy $\sum_\ell B_\ell = B$.

We will show that the iterations (8) converge to an achievable vector $\boldsymbol{B}$ under which the total delay threshold $\sum_\ell B_\ell$ is less than B.

We will utilize the ODE method ([28]–[30]) in order to analyze the discrete recursions (8). We will build some machinery in order to be able to utilize the ODE method. In a nutshell, the ODE method says that the discrete recursions (8) asymptotically track the ode,

$$\dot{B}_\ell = B_\ell(t) \left(\bar{\lambda}_\ell(\boldsymbol{B}(t)) - \sum_\ell \bar{\lambda}_\ell(\boldsymbol{B}(t)) B_\ell(t) \right), \forall \ell \in E. \tag{9}$$

A more precise statement is as follows.

Lemma 2: Define a continuous time process $\tilde{\boldsymbol{B}}(t), t \geq 0, t \in \mathbb{R}$ in the following manner. Define time instants $s(n), n = 0, 1, \ldots$ by setting $s(n) = \sum_{m=0}^{n-1} \gamma_m$. Now let

$$\tilde{\boldsymbol{B}}(s(n)) = \boldsymbol{B}(n), n = 0, 1, 2, \ldots.$$

Then use linear interpolation on each interval $[s(n), s(n+1)]$ to obtain values of $\tilde{\boldsymbol{B}}(t)$ for the time interval $[s(n), s(n+1)]$. Then for any $T > 0$,

$$\lim_{s\to\infty} \sup_{t\in[s,s+T]} \|\tilde{\boldsymbol{B}}(t) - \boldsymbol{B}^s(t)\|^2 = 0, \tag{10}$$

where $\boldsymbol{B}^s(t), t \geq s$ is a solution of (9) on $[s, \infty)$ with $\boldsymbol{B}^s(t) = \tilde{\boldsymbol{B}}(s)$. Thus, the recursions can be viewed as discrete analogue of the corresponding deterministic ode.

Proof: A detailed proof of (10) can be found in [28] Ch:5 Theorem 2.1. ∎

We note that the discrete-time process $\boldsymbol{B}(k), k = 0, 1, \ldots$ of (8) is deterministic. However, the recursions (8) do assume that the price vector $\bar{\lambda}(\boldsymbol{B}(k))$ is available in order to carry out the update.

It now follows from Lemma 2 that in order to study the convergence properties of the recursions (8), it suffices to analyze the detrministic ode (9).

We identify the above ode as the replicator equation [19], which has been utilized in evolutionary game theory. Next, we will make some assumptions regarding the replicator ode (9), and show that the trajectory of the ode has the desired properties under these assumptions. Let $\langle \cdot, \cdot \rangle$ denote the dot product of two vectors.

Definition 1: The ode (9) satisfies the monotonicity condition [31] if

$$\left\langle \boldsymbol{B} - \hat{\boldsymbol{B}}, \bar{\lambda}(\boldsymbol{B}) - \bar{\lambda}(\hat{\boldsymbol{B}}) \right\rangle < 0, \forall \boldsymbol{B} \neq \hat{\boldsymbol{B}} \in S. \tag{11}$$

We have,

Lemma 3: If the ode (9) satisfies the monotonicity condition, then for $\boldsymbol{B}(0)$ in the interior of S, $\boldsymbol{B}(t)$, the solution of the ode (9) satisfies $\boldsymbol{B}(t) \to \boldsymbol{B}^\star$, which is the unique equilibrium point of (9).

Lemmas 2 and 3 guarantee that the algorithm that starts with a suitable value of the delay budget vector $\boldsymbol{B}(0)$, and then tunes it according to the rule (8) converges to a feasible value of the delay budget vector $\boldsymbol{B}^\star$. Once the feasible $\boldsymbol{B}^\star$ has been obtained, the scheme (23) can be used in order to route packets. In summary,

Theorem 2: [Replicator Dynamics based $\boldsymbol{B}$ tuner] For the iterations (8), we have that $\boldsymbol{B}(k) \to \boldsymbol{B}^\star$, and hence the average queue lengths under the scheme described above are bounded by B.

Proof: It follows from Lemma 2 that the recusions (8) asymptotically track the ode (9) in the "$\gamma \to 0, k \to \infty$" limit (see Ch: 5 of [28], or Ch: 2 of [30]). Since the trajectory $\boldsymbol{B}(t)$ of the corresponding ode (9) satisfies $\boldsymbol{B}(t) \to \boldsymbol{B}^\star$, the proof follows. ∎

Remark 1: It must be noted that whether or not the delay bound B is achievable depends upon the network characteristics and traffic arrivals. Hence, if the bound B is too ambitious, it may not be achievable. In this case, the condition (11) will clearly not hold true, and the evolutionary algorithm will not be able to allocate the delay budgets appropriately. Thus, it is required that the network operator calculate an estimate of the value of the bound B that can be achieved, using knowledge of the arrival processes, underlay link capacities etc. Another way to choose an appropriate value of B can be to simulate the system performance under a "reasonably good policy", and set B equal to the average delays incurred under this policy.

Remark 2: We note the important role that the link-prices λ_ℓ play in order to "signal" to the tuner about the underlay network's characteristics. Thus, for example, a link ℓ might be strategically located, so that a large chunk of the network traffic necessarily needs to be routed through it, thus leading to a large value of average queue lengths under *any* routing policy π. Alternatively, its reliability might be low, which again leads to large value of queue lengths $Q_\ell(t)$. In either case, the quantity $B_\ell^\star$ would converge to a "reasonably large value", which in turn is enabled by high values of prices λ_ℓ during the course of iterations (8).

A. Tuning $\boldsymbol{B}$ using online learning

The scheme proposed in Theorem 2 needs to compute the prices $\bar{\lambda}(\boldsymbol{B}(k))$ in order to carry out the k-th update. One way to compute the $\bar{\lambda}(\boldsymbol{B}(k))$ is to apply Algorithm (23) with thresholds set to $\boldsymbol{B}(k)$, and wait for the link prices to converge to $\bar{\lambda}(\boldsymbol{B}(k))$. In order to speed up this naive scheme, we would like to carry out the B_ℓ updates in parallel with the iterations (23)-(24), an idea that is similar to the two timescale stochastic approximation scheme of Section VIII-A3. Thus, we propose the following scheme that comprises of three iterative update processes evolving simultaneously,

$$V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t)) \leftarrow \left\{V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t))\right\}(1-\alpha_t) + \alpha_t \left\{ \sum_{\ell\in R_f} \lambda_\ell Q_\ell^f(t) + \min_{\boldsymbol{u}^f} V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f) - V(\boldsymbol{q}_0, \boldsymbol{u}_0) \right\}, \tag{12}$$

$$\lambda_\ell(t+1) = \mathcal{M}\left\{\lambda_\ell(t) + \beta_t\left(\|\boldsymbol{Q}_\ell(t)\| - B_\ell\right)\right\}, \forall \ell \in E, \tag{13}$$

$$B_\ell(t+1) = \Gamma\left(B_\ell(t) + \gamma_t\left\{B_\ell(t)\left(\lambda_\ell(t) - \sum_\ell \lambda_\ell(t) B_\ell(t)\right)\right\}\right), \quad \forall \ell \in E, \tag{14}$$

where the step sizes satisfy $\beta_t = o(\alpha_t), \gamma_t = o(\beta_t)$. The condition $\gamma_t = o(\beta_t)$ will ensure that the $\boldsymbol{B}$ recursions view the network-wide link prices as having equilibriated for a fixed value of $\boldsymbol{B}$.

Next, we show that such a scheme, depicted in Algorithm 1, in which the $\boldsymbol{B}$ updates occur in parallel with the price cum Q-learning iterations, does indeed converge. The analysis is similar to the two time-scale stochastic approximation scheme discussed in the previous section. Its proof is provided in Appendix.

Theorem 3: [Three Time-scale POC Algorithm] For the iterations (8), almost surely, $\boldsymbol{B}(t) \to \boldsymbol{B}^\star$. Hence, the average

Algorithm 1 Price-based Overlay Controller (POC)

Fix sequences $\alpha_t, \beta_t, \gamma_t$ satisfying $\sum \alpha_t^2 < \infty, \sum \alpha_t = \infty, \sum \beta_t^2 < \infty, \sum \beta_t = \infty, \sum \gamma_t^2 < \infty, \sum \gamma_t = \infty$, and $\beta_t = o(\alpha_t), \gamma_t = o(\beta_t)$.
Initialize $\lambda(0) > 0$, $\boldsymbol{B}(0)$ such that $\sum_\ell B_\ell(0) = B$.
Perform the following iterations.

1) Each source node s_f updates its Q-values using,

$$V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t)) \leftarrow \left\{V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t))\right\}(1-\alpha_t) + \alpha_t \left\{ \sum_\ell \lambda_\ell Q_\ell^f(t) + \min_{\boldsymbol{u}^f} V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f) - V(\boldsymbol{q}_0, \boldsymbol{u}_0) \right\} \tag{15}$$

and the routing action implemented for flow f at time t is $\arg\max_{\boldsymbol{u}^f \in U_f} V(\boldsymbol{Q}^f(t), \boldsymbol{u}_f)$.

2) Overlay updates the link prices according to,

$$\lambda_\ell(t+1) = \mathcal{M}\left\{\lambda_\ell(t) + \beta_t\left(\|\boldsymbol{Q}_\ell(t)\| - B_\ell\right)\right\}, \forall \ell \in E. \tag{16}$$

3) Overlay adapts the link-level delay requirements according to,

$$B_\ell(t+1) = \Gamma\left(B_\ell(t) + \gamma_t\left\{B_\ell(t)\left(\lambda_\ell(t) - \sum_\ell \lambda_\ell(t) B_\ell(t)\right)\right\}\right), \tag{17}$$

$\forall \ell \in E$.

queue lengths suffered under the Price-based Overlay Controller (POC) described in Algorithm 1 are bounded by B.

V. Putting it all together: Three Layer POC

Algorithm 1 is composed of the following three layers from top to bottom:

1) *Replicator Dynamics* based B_ℓ tuner that observes the link prices λ_ℓ, and increases/decreases B_ℓ appropriately while keeping their sum equal to B.
2) *Sub-gradient descent* based λ_ℓ tuner that is provided the delay budget B_ℓ by the layer 1 described above, observes the queue lengths $\boldsymbol{Q}_\ell$ and updates the prices based on the mismatch between B_ℓ and $\bar{Q}_\ell$.
3) *Q-learning* based routing policy learner which is provided the link prices λ_ℓ, and learns to minimize the holding cost $\lambda_\ell \bar{Q}_\ell^f$. Its actions reflect the congestions on various links through the outcomes $\bar{Q}_\ell^f$, which in turn are used by layer 2 above.

The above three layers are thus intimately connected and coordinate amongst themselves in order to attain the goal of keeping the average queue lengths bounded by B. This is shown in Figure 2.

VI. Complexity of POC and possible Extensions

We discuss about the enormous simplification and reduced computational complexity of the POC as compared to the centralized optimal controller that has been obtained by solving the constrained MDP (1) using Linear Prgoramming approach. Firstly, the link prices λ_ℓ play the role of coordinating the actions of different flows. This allows us to reduce the problem of jointly scheduling F flows (1), to much simpler sub-problems that involve designing policy for only a single flow. As discussed earlier, the computational complexity scaled linearly with the number of flows, as compared with exponential in F.

The algorithms proposed until now assume that the overlay source nodes s_f know which underlay links route its packets. If these are unknown, then we can utilize Algorithm 1 on the overlay network comprising of tunnels. Thus, each overlay source node s_f maintains "virtual queues" $\hat{Q}_\tau^f(t)$ for each of its outgoing tunnel τ. Since link-level queue lengths are unknown, $\hat{Q}_\tau^f(t)$ are set to be equal to the "number of packets in flight" on tunnel τ, i.e., the number of packets that have been sent on tunnel τ, but have not reached their destination node by time t. The link prices and link-level average queue threshold requirements are now replaced by their tunnel counterparts $\lambda_\tau(t), B_\tau(t)$. Their iterations proceed in exactly the same manner as in Algorithm 1. The routing algorithm is much simpler, so that each source routes packets on its tunnel with the least price $\lambda_\tau(t)$. This yields the tunnel level POC, which is denoted POC-T (tunnel level POC) described in Algorithm 2.

Algorithm 2 Tunnel Level POC (POC-T)

Fix sequences $\alpha_t, \beta_t, \gamma_t$ satisfying conditions of Algorithm 1.
Initialize $\lambda(0) > \mathbf{0}$, $\boldsymbol{B}(0)$ such that $\sum_\tau \boldsymbol{B}_\tau(0) = B$.
Perform the following operations.

1) Each source node s_f updates the tunnel queues,

$$\hat{Q}_\tau^f(t) = \hat{Q}_\tau^f(t-1) + A_\tau^f(t) - D_\tau^f(t),$$

where $A_\tau^f(t)$ are the number of flow f packets that are routed to tunnel τ at time t, and $D_\tau^f(t)$ are the number of flow f packets that had been routed through tunnel τ, and reach their destination node at time t. The node s_f routes packets to the tunnel τ with the least holding cost $\lambda_\tau(t)\hat{Q}_\tau^f(t)$.

2) Tunnel prices are updated as,

$$\lambda_\tau(t+1) = \mathcal{M}\left\{\lambda_\tau(t) + \beta_t\left(\hat{Q}_\tau(t) - B_\tau\right)\right\}.$$

3) The tunnel-level average queue thresholds B_τ are updated according to,

$$B_\tau(t+1) = \Gamma\left(B_\tau(t) + \gamma_t\left\{B_\tau(t)\left(\lambda_\tau(t) - \sum_\tau \lambda_\tau(t) B_\tau(t)\right)\right\}\right).$$

Another important generalization is to consider the case where each flow f requires its average queue lengths $\|\bar{\boldsymbol{Q}}^f\|$ to be bounded by a threshold. The POC Algorithm can bemodified so that it now maintains B_ℓ^f, i.e., link-level queue

bounds for each flow f. Similarly, links maintain λ_ℓ^f, i.e., separate prices for each flow f and the routing algorithm for flow f seeks to minimize the holding cost $\sum_{\ell \in R_f} \lambda_\ell^f \bar{Q}_\ell^f$.

We can also consider the case where underlay implements a dynamic policy such as largest queue first, or backpressure policy [32]. The convergence of our proposed algorithms to the optimal policy, i.e., Theorems 4, 3 will continue to hold true in this set-up. Denote by π_{ul} the policy being implemented at the underlay. The proof of Theorem 4 will then be modified by using the fact that for a fixed value of price vector λ, the steady-state control policy applied to the network is given by $(\pi^\star(\lambda), \pi_{ul})$. Since the resulting controlled network still evolves on a finite state-space, and the policy $(\pi^\star(\lambda), \pi_{ul})$ is stationary, it is easily verified that the conclusions of Lemmas 6-8, Theorem 4, and the results of Section IV-A hold true.

Yet another possibility is to consider the case where the data packets have hard deadline constraints, i.e., they should reach their destination within a prespecified deadline in order for them to be counted towards "timely-throughput". Real-time traffic usually makes such stringent requirements on meeting end-to-end deadlines. A metric to judge the performance of scheduling policies for such traffic is the timely-throughput metric [33]. Timely throughput of a flow is the average number of packets per unit time-slot that reach their destination node within their deadline. The POC algorithm can be modified in order to maximize the cumulative timely-throughput of all the flows in the network. The packet-level router would then maximize the timely throughput minus the cost associated with using link bandwidth, which is priced at λ_ℓ units. Moreover, the topmost Replicator Dynamics based layer will not be required, and will be removed. The pricing updates will now be modified, so that they will try to meet the link capacities C_ℓ, and not the average queue delays at link ℓ. A detailed treatment for the case when the network is controllable can be found in [34].

In this paper we have allowed for control at only the source nodes. In practice, it may be desirable to control some intermediate nodes, since that would enable us to attain a lower value of average queue lengths. A modification to the bottommost layer will enable POC to handle such situations. Thus, each controllable node i will have to perform Q-learning updates (12) in order to "learn" its optimal routing policy. Thereafter, at each time t, it will make routing choices based on the instantaneous queue lengths $\boldsymbol{Q}^f(t)$, and the link prices λ by utilizing the optimal routing policy generated via Q-learning.

VII. Conclusions

We have analyzed optimal overlay routing problem in the system where the network operator has to satisfy a performance bound of average end-to-end delays less than D units. The problem is challenging because the overlay does not know the underlay characteristics. We proposed a simple decentralized algorithm based on a 3 layered adaptive control design. The algorithms could easily be tuned to operate under a vast multitude of information available about the underlay congestion. We also propose a heuristic scheme in case the flows do not know the links that are used to route their traffic. Simulations results show that the proposed schemes significantly outperform the existing policies.

References

[1] R. K. Sitaraman, M. Kasbekar, W. Lichtenstein, and M. Jain, "Overlay networks: An akamai perspective," *Advanced Content Delivery, Streaming, and Cloud Services*, pp. 305–328, 2014.

[2] L. L. Peterson and B. S. Davie, *Computer networks: a systems approach*. Elsevier, 2007.

[3] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, Dec 1992.

[4] A. L. Stolyar, "Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic," *The Annals of Applied Probability*, vol. 14, no. 1, pp. 1–53, 02 2004.

[5] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, *Resilient overlay networks*. ACM, 2001, vol. 35, no. 5.

[6] J. Han, D. Watson, and F. Jahanian, "Topology aware overlay networks," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 4. IEEE, 2005, pp. 2554–2565.

[7] N. M. Jones, G. S. Paschos, B. Shrader, and E. Modiano, "An overlay architecture for throughput optimal multipath routing," in *Proceedings of the 15th ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2014, pp. 73–82.

[8] G. S. Paschos and E. Modiano, "Throughput optimal routing in overlay networks," in *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*. IEEE, 2014, pp. 401–408.

[9] J. A. Boyan and M. L. Littman, "Packet routing in dynamically changing networks: A reinforcement learning approach," *Advances in neural information processing systems*, pp. 671–671, 1994.

[10] B. Awerbuch and R. D. Kleinberg, "Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches," in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, 2004, pp. 45–53.

[11] R. Gallager, "A minimum delay routing algorithm using distributed computation," *IEEE transactions on communications*, vol. 25, no. 1, pp. 73–85, 1977.

[12] J. Tsitsiklis and D. Bertsekas, "Distributed asynchronous optimal routing in data networks," *IEEE Transactions on Automatic Control*, vol. 31, no. 4, pp. 325–332, 1986.

[13] L. Bui, R. Srikant, and A. Stolyar, "Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 2936–2940.

[14] S. Asmussen, *Applied Probability and Queues*. Wiley, 1987.

[15] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *The Journal of the Operational Research Society*, vol. 49, no. 3, pp. pp. 237–252, 1998.

[16] K. J. Astrom and B. Wittenmark, *Adaptive Control*, ser. Dover Books on Electrical Engineering. Dover Publications, 2008.

[17] V. Borkar and P. Varaiya, "Adaptive control of Markov chains, I: Finite parameter set," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 953–957, Dec 1979.

[18] D. P. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, 2006.

[19] J. Weibull, "Evolutionary game theorymit press," *Cambridge, MA*, 1995.

[20] J.C. Gittins, K. Glazebrook and R. Weber, *Multi-armed Bandit Allocation Indices*. John Wiley & Sons, 2011.

[21] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.

[22] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice hall Englewood Cliffs, NJ, 1989, vol. 23.

[23] E. Altman, *Constrained Markov Decision Processes*. Chapman and Hall/CRC, March 1999.

[24] D. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Athena Scientific, 2001, vol. 1 and 2.

[25] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st ed. New York, NY, USA: John Wiley & Sons, Inc., 1994.
[26] V. S.Borkar, "A convex analytic approach to Markov decision processes," *Probability Theory and Related Fields*, vol. 78, no. 4, pp. 583–602, 1988.
[27] D. P. Bertsekas, A. E. Ozdaglar, and A. Nedic, *Convex analysis and optimization*, ser. Athena scientific optimization and computation series. Belmont (Mass.): Athena Scientific, 2003.
[28] H. J. Kushner and G. Yin, *Stochastic Approximation Algorithms and Applications*. New York: Springer Verlag, 1997.
[29] V. S. Borkar and S. P. Meyn, "The o.d.e. method for convergence of stochastic approximation and reinforcement learning," *SIAM Journal on Control and Optimization*, vol. 38, no. 2, pp. 447–469, 2000.
[30] V. S. Borkar, *Stochastic Approximation : A Dynamical Systems Viewpoint*. Cambridge: Cambridge University Press New Delhi, 2008.
[31] H. Smith, *Monotone dynamical systems*. Providence, RI: American Mathematical Society, 1995.
[32] "Technical report," https://www.dropbox.com/s/v5d5svfpypfdieq/technical%20report.pdf?dl=0.
[33] I-Hong Hou and V.S. Borkar and P.R. Kumar, "A Theory of QoS for Wireless," in *IEEE INFOCOM 2009*, April 2009, pp. 486–494.
[34] R. Singh and P. Kumar, "Throughput optimal decentralized scheduling of multi-hop networks with end-to-end deadline constraints: Unreliable links," *arXiv preprint arXiv:1606.01608*, 2016.
[35] J. Abounadi, D. Bertsekas, and V. S. Borkar, "Learning algorithms for markov decision processes with average cost," *SIAM Journal on Control and Optimization*, vol. 40, no. 3, pp. 681–698, 2001.
[36] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998. [Online]. Available: http://www.cs.ualberta.ca/\%7Esutton/book/ebook/the-book.html
[37] J. N. Tsitsiklis, "Asynchronous stochastic approximation and q-learning," *Machine Learning*, vol. 16, no. 3, pp. 185–202, 1994.
[38] T. Jaakkola, M. I. Jordan, and S. P. Singh, "On the convergence of stochastic iterative dynamic programming algorithms," *Neural computation*, vol. 6, no. 6, pp. 1185–1201, 1994.
[39] H. J. Kushner and D. S. Clark, *Stochastic approximation methods for constrained and unconstrained systems*. Springer Science & Business Media, 2012, vol. 26.
[40] V. R. Konda and J. N. Tsitsiklis, "Actor-critic algorithms." in *NIPS*, vol. 13, 1999, pp. 1008–1014.
[41] V. S. Borkar, "An actor-critic algorithm for constrained markov decision processes," *Systems & control letters*, vol. 54, no. 3, pp. 207–213, 2005.
[42] W. F. Ames and B. Pachpatte, *Inequalities for differential and integral equations*. Academic press, 1997, vol. 197.
[43] K. C. Border, *Fixed point theorems with applications to economics and game theory*. Cambridge university press, 1989.

VIII. Appendix

A. Using Multiple Time-scale based Online Learning methods

1) Employing Online Learning to Obtain $\pi^\star_f(\lambda)$: We can utilize online learning methods such as Relative Value Iteration (RVI) *Q-learning* [35] in order to obtain $\pi^\star_f(\lambda)$, i.e., the policy that minimizes the holding cost $\sum_\ell \lambda_\ell \bar{Q}^f_\ell$ for flow f. Each overlay source node s_f maintains the *Q-values* [35], [36] $V(\boldsymbol{q}^f, \boldsymbol{u}^f)$ which corresponds to the "long-term" utility obtained by taking the routing action $\boldsymbol{u}^f$ when the flow f queue vector is $\boldsymbol{q}^f$. These Q-values are updated using online data as follows,

$$V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t)) \leftarrow \left\{V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t))\right\}(1-\alpha_t) + \alpha_t \left\{ \sum_{\ell \in R_f} \lambda_\ell Q^f_\ell(t) + \min_{\boldsymbol{u}^f} V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f) - V(\boldsymbol{q}_0, \boldsymbol{u}_0) \right\}, \tag{18}$$

where $\alpha_t = \frac{1}{t}$, and $(\boldsymbol{q}_0, \boldsymbol{u}_0)$ is a fixed state-action pair. The routing action implemented at time t corresponds to the action $\boldsymbol{u}^f$ that minimizes the quantity $V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f)$. Convergence of the Q-learning scheme to the optimal policy is well-understood, and detailed convergence proofs can be found in [28], [30], [37], [38]. We however state this as a lemma since it will be required in convergence analysis of later sections, when we will combine the Q-learning iterations with other iterative algorithms.

Lemma 4: The asymptotic properties of the stochastic iterations (18) can be derived by studying the following ordinary differential equation (ode),

$$\dot{V} = TV - eV(\boldsymbol{q}_0, \boldsymbol{u}_0), \tag{19}$$

where T is the Bellman operator, i.e., $TV(\boldsymbol{q}, \boldsymbol{u}) = \|\boldsymbol{q}\| + \mathbb{E}\{\min_u V(\tilde{\boldsymbol{q}}, \boldsymbol{u})\}$, and e is vector comprising of all ones. Since the mapping $V \mapsto TV - V$ is contractive, the ode (19), and hence the Q-learning iterations (18) converge, thus yielding the optimal routing policy.

Proof: See [35]. ∎

We note that the Q-learning algorithm as discussed above learns the optimal policy, i.e., a mapping from the instantaneous queue lengths $\boldsymbol{Q}^f(t)$ to a routing action $\boldsymbol{U}^f(t)$. It does not require the knowledge of the dynamics of queue lengths evolution in the network, and this is precisely the reason why we employ Q-learning, because the queue dynamics are unknown to the policy at the overlay.

2) Obtaining $\lambda^\star$ using Subgradient Descent Method: Next, we address the issue of obtaining the optimal value of vector of link prices $\lambda^\star$. Let us, for the time being, assume that for a given value of link prices λ, the policy $\pi^\star(\lambda)$ has been obtained. This could be achieved, for example, by performing the Q-learning iterations until convergence. Since,

$$D(\lambda) = \mathcal{L}(\pi^\star(\lambda), \lambda), \tag{20}$$

and from (5),

$$\frac{\partial \mathcal{L}}{\partial \lambda_\ell} = \|\bar{\boldsymbol{Q}}_\ell\| - B_\ell, \quad \forall \ell \in E,$$

we can use the gradient descent method [27] to compute $\lambda^\star$ via the following iterations,

$$\lambda_\ell(k+1) = \lambda_\ell(k) + \beta_k \left(\|\bar{\boldsymbol{Q}}_\ell\| - B_\ell\right), \forall \ell \in E, \tag{21}$$

where $\|\bar{\boldsymbol{Q}}_\ell\|$ is the steady state average queue lengths at link ℓ under the application of $\pi^\star(\lambda(k))$, while β_k can be taken to be $\frac{1}{k}$. We state the convergence result since it will be utilized later.

Lemma 5: The following ode

$$\dot{\lambda} = \nabla_\lambda D(\lambda) \tag{22}$$

converges to $\lambda^\star$ the solution of the dual problem (20), and hence the iterations (21) converge to $\lambda^\star$. Since the price iterations (21) track the ode (22) in the $\beta_k \to 0, k \to \infty$, they too converge to optimal price $\lambda^\star$.

Proof: In order to analyze the properties of the ode (22), let us take the value of the dual function D. We then have,

$$\frac{d}{dt} D(\lambda(t)) = -\|\nabla_\lambda D(\lambda(t))\|^2 \leq 0.$$

Since D is bounded from below, this proves the claim made above.

The connection between the ode and its discrete counterpart is made using Kushner-Clarke theorem, see [30] Ch: 2 or [39] Ch:2, or [28] Ch:5. ∎

We note that here the variable k is used to index the price iterations, and must not be confused with the time index t. The price update $\lambda(k) \to \lambda(k+1)$ occurs only when the Q-learning iterations (18) with prices set to $\lambda(k)$ have converged.

3) Two Time Scale Stochastic Approximation: As discussed above, successive price updates under the scheme (21) need to wait for the Q-learning iterations to converge to the policy $\pi^\star_f(\lambda(k))$. Thus, a learning scheme that combines the iterations of Sections VIII-A1 and VIII-A2 would suffer from slow convergence. An alternative would be to perform both the iterations simultaneously, though the price updates be carried out on a slower time-scale than that of Q-learning. This would enable the price iterations to view the Q-learning iterations as having converged, and hence the time-averaged queue lengths equilibriated to their steady-state values. Similarly, the Q-learning iterations (18) view the link prices λ as static. Thus, we propose the following scheme which combines both these iterations,

$$V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t)) \leftarrow \left\{V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t))\right\}(1-\alpha_t) + \alpha_t \left\{ \sum_{\ell \in R_f} \lambda_\ell Q^f_\ell(t) + \min_{\boldsymbol{u}^f} V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f) - V(\boldsymbol{q}_0, \boldsymbol{u}_0) \right\} \quad (23)$$

$$\lambda_\ell(t+1) = \mathcal{M}\left\{\lambda_\ell(t) + \beta_t\left(\|\boldsymbol{Q}_\ell(t)\| - B_\ell\right)\right\}, \forall \ell \in E, \quad (24)$$

where $\mathcal{M}(\cdot)$ is the operator that projects the price onto the compact set $[0, K]$ for some sufficiently large $K > 0$[5], and the sequence β_t satisfies $\beta_t = o(\alpha_t)$, i.e., $\lim_{t\to\infty} \frac{\beta_t}{\alpha_t} = 0$. For example, one could let $\alpha_t = \frac{1}{t}, \beta_t = \frac{1}{t \log t}$.

The above two-time scale stochastic approximation algorithm can be shown to converge, thereby yielding optimal prices $\lambda^\star$, and also the optimal policy $\pi^\star(\lambda^\star)$. We will perform an analysis of algorithm (23) using the two time-scale ODE method [28], [30].

Let us denote by $\boldsymbol{V}_\lambda$ the vector comprising of the converged values of "cost-to-go" function $V(\cdot)$ when the RVI Q-learning iterations (18) are performed with prices set to λ. Since the price iterations are performed on a slower time-scale than Q-learning, we would expect that asymptotically the routing policy being implemented under the Q-learning scheme is given by $\pi^\star(\lambda(t))$. The following result formalizes this.

Lemma 6: For the scheme (23)-(24) that adjusts the Q-values and link prices simultaneously, we have,

$$(\boldsymbol{V}(t), \lambda(t)) \to \{(\boldsymbol{V}_\lambda, \lambda)\} \text{ almost surely (a.s.), and }, \quad (25)$$

$$\pi(t) \to \pi(\boldsymbol{V}_\lambda) \text{ a.s. }. \quad (26)$$

Proof: We will provide a proof sketch. A detailed proof is along the lines of Section 6 of [40], or Lemma 4.1 of [41].

[5]Such a projection is required in order to show that the "noisy"/stochastic iterations converge.

The stochastic recursions (23) for prices can be re-written as follows, so that they are performed on the same time-scale α_t,

$$V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t)) \leftarrow \left\{V(\boldsymbol{Q}^f(t), \boldsymbol{U}^f(t))\right\}(1-\alpha_t) + \alpha_t \left\{ \sum_{\ell \in R_f} \lambda_\ell Q^f_\ell(t) + \min_{\boldsymbol{u}^f} V(\boldsymbol{Q}^f(t+1), \boldsymbol{u}^f) - V(\boldsymbol{q}_0, \boldsymbol{u}_0) \right\}, \quad (27)$$

$$\lambda_\ell(t+1) = \mathcal{M}\left\{\lambda_\ell(t) + \alpha_t\left[\frac{\beta_t}{\alpha_t}\left(\|\bar{\boldsymbol{Q}}_\ell(t)\| - B_\ell\right)\right]\right\}. \quad (28)$$

Since $\frac{\beta_t}{\alpha_t} \to 0$, it then follows that the following system of odes can be analyzed in order to infer asymptotic properties of (27),

$$\dot{V} = TV - eV,$$
$$\dot{\lambda}_\ell = 0, \forall \ell \in E.$$

The result thus follows. ∎

Let us now analyze the price iterations on the slower time-scale β_t. Letting

$$g_\ell(\boldsymbol{Q}(t-1), \boldsymbol{U}(t-1)) = \mathbb{E}\left\{\|\boldsymbol{Q}_\ell(t)\| - B_\ell | \boldsymbol{Q}(t-1), \boldsymbol{U}(t-1)\right\},$$

be the "mean drift" of the system at time t under the action $\boldsymbol{U}(t)$, we have

$$\lambda_\ell(t+1) = \lambda_\ell(t) + \beta_t\left(g(\boldsymbol{Q}(t), \boldsymbol{U}(t)) + \psi_\ell(t)\right), \forall \ell \in E$$

where $\psi_\ell(t) := (\|\boldsymbol{Q}_\ell(t)\| - B_\ell) - g_\ell(\boldsymbol{Q}(t-1), \boldsymbol{U}(t-1))$ is the martingale difference "noise". Lemma 27 suggests that the behaviour of price iterations should not change much if the routing policy $\pi^\star(\lambda(t))$ was utilized, rather than the control action $\boldsymbol{U}(t)$ chosen according to the RVI Q-learning that is being applied under our scheme. Thus, the price iterations can be re-written as

$$\lambda_\ell(t+1) = \lambda_\ell(t) + \beta_t g_\ell(\boldsymbol{Q}(t), \pi^\star(\lambda(t))) + \beta_t\left(g_\ell(\boldsymbol{Q}(t-1), \boldsymbol{U}(t-1)) - g_\ell(\boldsymbol{Q}(t), \pi^\star(\lambda(t)))\right) + \beta_t \psi_\ell(t), \forall \ell \in E. \quad (29)$$

The recursions

$$\lambda_\ell(t+1) = \lambda_\ell(t) + \beta_t g(\boldsymbol{Q}(t), \pi^\star(\lambda(t))), \forall \ell \in E \quad (30)$$

fall into the category of "stochastic approximation with Markov noise" that is typically considered in adaptive control. This set-up consists of a system whose dynamics are unknown, and at each time t, the control action is chosen according to a control law that is parametrized by a parameter $\lambda(t)$. Moreover, the parameter $\lambda(t)$ is also tuned based on the cost incurred by the system. Recursions (30) correspond to the set-up in which the unknown underlay network is controlled using the control law $\pi^\star(\lambda(t))$, which is parametrized by the parameter $\lambda(t)$. Such kind of stochastic recursive algorithms have been analyzed in detail in literature, and a detailed discussion can be found in [28], [30]. We state the convergence result.

Lemma 7: For the recursions (30), we have that

$$\lambda(t) \to \lambda^\star,$$

where $\lambda^\star$ is the solution to the dual problem (7).

Proof: See Theorem 7 in Ch: 6 of [30]. ∎

Let us now get back to analyzing the original price recursions (29). Using (29), the price evolution over multiple time-steps can be written as,

$$\lambda(m+n) = \lambda(m) + \sum_{k=m}^{m+n} \beta_k g_\ell(\boldsymbol{Q}(k), \pi^\star(\lambda(k))), + \sum_{k=m}^{m+n} \beta_k \left[g_\ell(\boldsymbol{Q}(k), \boldsymbol{U}(k)) - g_\ell(\boldsymbol{Q}(k), \pi^\star(\lambda(k))) \right] + \sum_{k=m}^{m+n} \beta_k \psi_\ell(k).$$

We will be interested in showing that the contribution of the second and third term in the r.h.s. is asymptotically neglegible. Now, the term $\sum_{k=m}^{m+n} \beta_k \psi(k)$ is the summation of martingale noise that "drives" the system, and is bounded by $\sup_{i>m} \sum_k \beta_k \psi(k)$, a quantity which is $O(\frac{1}{m})$. It follows from the second assertion in Lemma 6, that the term $\sum_{k=m}^{m+n} \beta_k \left[g(\boldsymbol{Q}(k), \boldsymbol{U}(k)) - g(\boldsymbol{Q}(k), \pi^\star(\lambda(k))) \right]$ vanishes as $m \to \infty$. In summary,

Lemma 8: For the price iterations (29), we have that,

$$\lim_{m\to\infty} \sum_{k=m}^{m+n} \beta_k \left[g(\boldsymbol{Q}(k), \boldsymbol{U}(k)) - g(\boldsymbol{Q}(k), \pi^\star(\lambda(k))) \right] = 0, \ a.s. \tag{31}$$

$$\text{and } \lim_{m\to\infty} \sum_{k=m}^{m+n} \beta_k \psi(k) = 0 \ a.s.. \tag{32}$$

Having built the machinery to analyze the two time-scale scheme (23)-(24), we now show its the converegence.

Theorem 4: [Link-Level Queue Control] Consider the iterative scheme (23)-(24) that performs the Q-learning (18) and the price iterations (21) simultaneously, though on different time-scales. The scheme converges, yields the optimal link prices $\lambda^\star$, and the optimal routing policy $\pi^\star(\lambda^\star)$. Thus, it solves the CMDP (3), i.e., it keeps the average queue lengths on each link ℓ bounded by the threshold B_ℓ.

Proof: We would like to show that the price iterations in (23)-(24) can be well approximated by the convergent ode (22). Thus, we begin by embedding the $(\boldsymbol{V}, \lambda)$ recursions into a continuous trajectory. Let $s(0) = 0$ and $s(k) = \sum_{i=0}^{k-1} \beta_k$. Now consider the function $\tilde{\lambda}$ defined as follows: set $\tilde{\lambda}(s(k)) = \lambda(k)$, and on each interval of the type $[s(k), s(k+1)]$ obtain the value of $\tilde{\lambda}(t)$ by using linear interpolation on the values $\tilde{\lambda}(s(k-1)), \tilde{\lambda}(s(k))$. The function $\tilde{\lambda}(t)$ thus obtained is continuous and piece-wise linear. Also define $[t]' := \max\{s(n) : s(n) \leq t\}$. For $s > 0$, define by $\lambda^s(t), t \geq s$, a trajectory of the ode (22) which satisfies $\lambda^s(s) = \tilde{\lambda}(s)$.

In order to be able to that the ode (22) does indeed approximate the price recursions in (23), it suffices to show that the following is true,

$$\lim_{s\to\infty} \sup_{t\in[s,s+T]} \|\tilde{\lambda}(t) - \lambda^s(t)\| = 0 \text{ almost surely.} \tag{33}$$

From the alternative formulation (29) of the price iterations, we have that,

$$\tilde{\lambda}(s(n+m)) = \tilde{\lambda}(s(n)) + \int_{s(n)}^{s(n+m)} \nabla_\lambda D(\tilde{\lambda}(t)) \mathrm{d}t + \int_{s(n)}^{s(n+m)} \nabla_\lambda D(\tilde{\lambda}([t]')) - \nabla_\lambda D(\tilde{\lambda}(t)) \mathrm{d}t + \sum_{k=1}^{m-1} \beta_{n+k} \left(\|Q_\ell(n+k)\| - \|Q_\ell(n+k)\|^{\pi^\star(\lambda(n+k))} \right) + \sum_{k=m}^{m+n} \beta_k \psi_k.$$

Denote the "error terms" in the above by $\Delta_1 := \int_{s(n)}^{s(n+m)} \nabla_\lambda D(\tilde{\lambda}([t]')) - \nabla_\lambda D(\tilde{\lambda}(t)) \mathrm{d}t$, $\Delta_2 = \sum_{k=1}^{m-1} \beta_{n+k} \left(\|Q_\ell(n+k)\| - \|Q_\ell(n+k)\|^{\pi^\star(\lambda(n+k))} \right)$, and $\Delta_3 = \sum_{k=m}^{m+n} \beta_k \psi_k$. Now the term Δ_1 vanishes because the function $\nabla_\lambda D(\lambda)$ (6) is continuous, while it follows from Lemma 8, that the terms Δ_2, Δ_3 vanish as $m \to \infty$.

Using Gronwall's inequality [42], we have that

$$\sup_{t\in[s,s+T]} \|\tilde{\lambda}(t) - \lambda^s(t)\| \leq K_T \left(\Delta_1 + \Delta_2 + \Delta_3 \right),$$

and hence the relation (33) holds true. Now using Theorem 2 in Chapter 2 of [30], it follows that $\lambda(t) \to \lambda^\star$ almost surely. The covergence of the policy, i.e., $\boldsymbol{U}(t) \to \pi^\star(\lambda^\star)$ follows from Lemma 6. ∎

Remark 3: Since the underlay links are uncontrollable, the link price iterations can be implemented at the overlay. The overlay however still needs to know the queue lengths $\boldsymbol{Q}(t)$ in order to implement the scheme (23).

Remark 4: Throughout this section, it was assumed that the link-level threshold $\boldsymbol{B}$ can be attained under some policy π. Section VIII-C discusses the case of infeasible $\boldsymbol{B}$.

B. Advantages of Flow-Level Decomposition

i) *Computational Advantages*: Let us denote the action space corresponding to the routing decisions for flow f by U_f, and by N_f the number of links which route flow f's packets. Then, since an overlay routing specifies an action for each possible value of network state, the number of *centralized* routing policies is given by $\prod_{f=1}^{F} |U_f|^{B^{N_f}}$. On the other hand, the number of decentralized policies is $\sum_{f=1}^{F} |U_f|^{B^{N_f}}$. Thus, the search for an optimal policy in the class of decentralized policies is computationally less expensive as compared to centralized policies.

ii) *Reduced Communication Overheads*: In order to implement a decentralized policy, each source node s_f needs to know only its own queue lengths $\boldsymbol{Q}^f(t)$. This eliminates the need to share the network queue lengths $\boldsymbol{Q}(t)$ globally amongst all the overlay nodes, thus saving significant communication overheads.

iii) *Parallel Distributed Computation* As we will show now, the optimal policy can be solved in a parallel and distributed fashion [22], in which each source node s_f

performs a search for flow f's optimal policy using online learning techniques.

C. Analysis of scheme (23)-(24) for the case of infeasible $\boldsymbol{B}$

If the delay budget $\boldsymbol{B}$ is achievable, then it follows from Theorem 4 that the algorithm converges to the optimal policy $\pi^\star$. This section will derive the asymptotic properties of price iteration when $\boldsymbol{B}$ is infeasible.

Recall that the price iterations were carried out in order to solve the following dual problem

$$\max_{\lambda} D(\lambda).$$

Since the dual function $D(\lambda)$ is concave, and the iterations were basically gradient ascent steps, they converged to the optimal price vector $\lambda^\star$. However, when the primal problem (3) is infeasible, the dual function $D(\lambda)$, though still concave, is unbounded. An application of the gradient descent method would make the prices $\lambda(t)$ unbounded. However since at each time t we project the prices onto a compact set, the price iterations solve the following constrained optimization problem,

$$\max_{\lambda:0\leq\lambda_\ell\leq K} D(\lambda),$$

where we note that the quantity K corresponds to the bound on prices λ_ℓ that is chosen by the scheme (23)-(24). The convergence properties of the algorithm for infeasible $\boldsymbol{B}$ then follows by combining [28] Ch: 5, with an analysis similar to that performed in Theorem 4. The key difference now will be an additional "forcing term" that appears when the iterates hit the "upper threshold" value of K. This additional forcing term keeps the iterates bounded.

Lemma 9: Consider the underlay network operating under the scheme (23), or equivalently the algorithm 1 implemented under a fixed value of $\boldsymbol{B}$ under which the problem (3) is infeasible. The iterates $\lambda(t)$ converge to a fixed point of the following ode,

$$\dot{\lambda} = \nabla_\lambda D(\lambda) + z(\lambda),$$

where $z(\lambda)$ is a "reflection term" that keeps the prices $\lambda(t)$ iterates bounded.

Corollary 1: Combining Theorem 4 and 9, we have that when algorithm 1 is utilized with the delay budget allocation kept fixed at $\boldsymbol{B}$, then the prices $\lambda(t) \to \bar{\lambda}(\boldsymbol{B})$.

D. Proof of Lemma 8

Proof: Let us show that there exists a unique $\boldsymbol{B}^\star \in S$ such that $\boldsymbol{B}^\star$ maximizes $\boldsymbol{B} \to (\boldsymbol{B})^\intercal \bar{\lambda}(\boldsymbol{B}^\star)$. It is easily verified that the set-valued map which maps $\boldsymbol{B} \in S$ to the set $\{\boldsymbol{x} \in S : \boldsymbol{x} \in \arg\max \boldsymbol{x}^\intercal \bar{\lambda}(\boldsymbol{B})\}$ is nonempty compact convex and uppersemicontinuous. It then follows from the Kakutani fixed point theorem [43] that there exists a $\boldsymbol{B}^\star$ which maximizes the function $\boldsymbol{x} \to \boldsymbol{x}^\intercal \bar{\lambda}(\boldsymbol{B}^\star)$ over the simplex S. Now assume that there exists another $\hat{\boldsymbol{B}} \neq \boldsymbol{B}^\star$ belonging to S, such that $\hat{\boldsymbol{B}}$ maximizes the function $\boldsymbol{x} \to \boldsymbol{x}^\intercal \bar{\lambda}(\hat{\boldsymbol{B}})$. This gives us,

$$\begin{aligned}\left\langle \boldsymbol{B}^\star - \hat{\boldsymbol{B}}, \bar{\lambda}(\boldsymbol{B}^\star) - \bar{\lambda}(\hat{\boldsymbol{B}}) \right\rangle &= (\boldsymbol{B}^\star)^\intercal \bar{\lambda}(\boldsymbol{B}^\star) \\ &- (\boldsymbol{B}^\star)^\intercal \bar{\lambda}(\hat{\boldsymbol{B}}) - (\hat{\boldsymbol{B}})^\intercal \bar{\lambda}(\boldsymbol{B}^\star) + (\hat{\boldsymbol{B}})^\intercal \bar{\lambda}(\hat{\boldsymbol{B}}) \\ &\geq 0,\end{aligned}$$

which contradicts our monotonicity assumption (11), and hence $\hat{\boldsymbol{B}} = \boldsymbol{B}^\star$.

Now we show that for $\boldsymbol{B}(0)$ in the interior of the set S, $\boldsymbol{B}(t) \to \boldsymbol{B}^\star$. Consider the Lyapunov function,

$$V(\boldsymbol{B}) := \sum_\ell B_\ell^\star \ln\left(\frac{B_\ell^\star}{B_\ell}\right), B \in S.$$

It follows from Jensen's inequality, that $V(\boldsymbol{B}) \geq 0$, and is 0 if and only if $\boldsymbol{B} = \boldsymbol{B}^\star$. Now consider,

$$\begin{aligned}\frac{\mathrm{d}V(\boldsymbol{B}(t))}{\mathrm{d}t} &= -\sum_\ell B_\ell^\star \left(\frac{\dot{B}_\ell(t)}{B_\ell(t)}\right) \\ &= (\boldsymbol{B}(t) - \boldsymbol{B}^\star)^\intercal \bar{\lambda}(\boldsymbol{B}(t)) \\ &\leq (\boldsymbol{B}(t) - \boldsymbol{B}^\star)^\intercal \bar{\lambda}(\boldsymbol{B}(t)) - (\boldsymbol{B}(t) - \boldsymbol{B}^\star)^\intercal \bar{\lambda}(\boldsymbol{B}^\star) \\ &= (\boldsymbol{B}(t) - \boldsymbol{B}^\star)^\intercal \left(\bar{\lambda}(\boldsymbol{B}(t)) - \bar{\lambda}(\boldsymbol{B}^\star)\right) \\ &< 0,\end{aligned}$$

for $\boldsymbol{B}(t) \neq \boldsymbol{B}^\star$. The first equality follows from (9), and first inequality is true because $\boldsymbol{B}^\star$ maximizes $(\boldsymbol{B}^\star)^\intercal \bar{\lambda}(\boldsymbol{B}^\star)$. The last inequality follows from the monotonicity assumption (11). The claim $\boldsymbol{B}(t) \to \boldsymbol{B}^\star$ follows. ∎

E. Proof of Theorem 3

As the first step in our analysis of Algorithm 1, we would like to show that since the price iterations are performed on a faster time-scale, the $\boldsymbol{B}$ tuner views them as averaged out, and hence the $\lambda(t)$ in the equation for $\boldsymbol{B}$ update can be approximated by $\bar{\lambda}(\boldsymbol{B}(t))$.

Lemma 10: For the iterates $\lambda(t), \boldsymbol{B}(t)$ in Algorithm 1 we have that $(\lambda(t), \boldsymbol{B}(t)) \to \{(\bar{\lambda}(\boldsymbol{B}), \boldsymbol{B} : \boldsymbol{B} \in S)\}$.

Proof: Since $\gamma_t = o(\beta_t)$, the analysis of the price and Q-learning iterations in Algorithm 1 can be carried out while keeping $\boldsymbol{B}$ as a constant. The result then follows from Corollary 1. ∎

We re-write the $\boldsymbol{B}$ iterations as follows,

$$B_\ell(t+1) - B_\ell(t) = \gamma_t \left\{ B_\ell(t) \left(\lambda_\ell(t) - \sum_\ell \lambda_\ell(t) B_\ell(t) \right) \right\}, \tag{34}$$

$$= \gamma_t \left\{ B_\ell(t) \left(\bar{\lambda}_\ell(\boldsymbol{B}(t)) - \sum_\ell \bar{\lambda}_\ell(\boldsymbol{B}) B_\ell(t) \right) \right\} + \gamma_t \Delta(t), \tag{35}$$

where,

$$\Delta(t) = B_\ell(t) \left(\lambda_\ell(t) - \bar{\lambda}_\ell(\boldsymbol{B}(t)) - \sum_\ell (\lambda_\ell(t) - \bar{\lambda}_\ell(\boldsymbol{B})) B_\ell(t) \right)$$

is the error term due to approximating the "true" price $\lambda(t)$ by its stationary value $\bar{\lambda}(\boldsymbol{B}(t))$. Define the piecewise-linear, continuous process $\tilde{\boldsymbol{B}}(t)$. We are interested in showing that asymptically the solutions of the ode (9) approximate well the values $\boldsymbol{B}(t)$ of the recursions in Algorithm 1. Thus,

Lemma 11: For a fixed $T > 0$ we have that

$$\lim_{s\to\infty} \sup_{t\in[s,s+T]} \|\tilde{\boldsymbol{B}}(t) - \boldsymbol{B}^s(t)\| = 0 \text{ almost surely.} \quad (36)$$

Proof: In the below, we let $f(\boldsymbol{B},\lambda) = B_\ell\left(\lambda_\ell(\boldsymbol{B}) - \sum_\ell \lambda_\ell(\boldsymbol{B})B_\ell\right)$. Consider the following,

$$\tilde{B}_\ell(s(n+m)) = \tilde{B}_\ell(s(n)) + \int_{s(n)}^{s(n+m)} f(\boldsymbol{B}(t), \bar{\lambda}(\boldsymbol{B}(t)))\mathrm{d}t$$
$$+ \left(\int_{s(n)}^{s(n+m)} f(\boldsymbol{B}([t]'), \bar{\lambda}(\boldsymbol{B}[t]')) - f(\boldsymbol{B}(t), \bar{\lambda}(\boldsymbol{B}(t)))\mathrm{d}t\right)$$
$$+ \sum_{k=m}^{m+n} \gamma_k f(\boldsymbol{B}(k), \lambda(k)) - f(\boldsymbol{B}(k), \bar{\lambda}(\boldsymbol{B}(k)))$$

Denote the "discretization error" by $\Delta_4 := \left(\int_{s(n)}^{s(n+m)} f(\boldsymbol{B}([t]'), \bar{\lambda}(\boldsymbol{B}[t]')) - f(\boldsymbol{B}(t), \bar{\lambda}(\boldsymbol{B}(t)))\mathrm{d}t\right)$. And let $\Delta_5 := \sum_{k=m}^{m+n} \gamma_k f(\boldsymbol{B}(k), \lambda(k)) - f(\boldsymbol{B}(k), \bar{\lambda}(\boldsymbol{B}(k)))$. We have $\Delta_4 \to 0$ as $s \to \infty$ since the function f is continuous. Since from Lemma 10 we have $\lambda(k) \to \bar{\lambda}(\boldsymbol{B}(k))$, and the function f is continuous, we deduce that $\Delta_5 \to 0$ as $s \to \infty$. ■

Proof of Theorem 3: The claim follows by combining Lemma 11 and the fact that the iterates $\boldsymbol{B}(t)$ are bounded, with Theorem 2 in Ch:2 of [30]. ■